Distributed Training of Deep Neural Networks with Spark: The MareNostrum Experience

Leonel Cruz[a], Ruben Tous[a,**], Beatriz Otero[a]

[a]*Universitat Politècnica de Catalunya, Barcelona, Spain*

ABSTRACT

Deployment of a distributed deep learning technology stack on a large parallel system is a very complex process, involving the integration and configuration of several layers of both, general-purpose and custom software. The details of such kind of deployments are rarely described in the literature. This paper presents the experiences observed during the deployment of a technology stack to enable deep learning workloads on MareNostrum, a petascale supercomputer. The components of a layered architecture, based on the usage of Apache Spark, are described and the performance and scalability of the resulting system is evaluated. This is followed by a discussion about the impact of different configurations including parallelism, storage and networking alternatives, and other aspects related to the execution of deep learning workloads on a traditional HPC setup. The derived conclusions should be useful to guide similarly complex deployments in the future.

1. Introduction

Over the past several years, deep neural networks (DNNs) have proven to be an incredibly effective tool for a variety of problems, from computer vision, speech recognition or natural language processing. Their number of parameters and complexity, and the size of the training datasets, have quickly grown, leading to be a first-class workload for HPC (High-Performance Computing) infrastructures. However, enabling deep learning workloads on a large parallel system is a very complex process, involving the integration and configuration of several layers of both, general-purpose and custom software. The details of such kind of deployments are rarely described in the literature. This paper presents the experiences observed during the deployment of a technology stack to enable deep learning workloads on a real-world, petascale, HPC setup, the MareNostrum supercomputer.

The goal of the deployment is to be able to take profit of the computation resources provided by MareNostrum (almost 50K cores and more than 100TB of aggregated RAM) for training DNNs. Nowadays, the usage of GPUs has proven to be the more efficient alternative to train neural networks, speeding up common operations such as large matrix computations (Lee et al. (2010); Fujimoto (2008)). As their price, performance and energy efficiency improves, GPUs are gaining ground in HPC (both in special-purpose systems and in hybrid general-purpose supercomputers). However, there are still many systems, such as MareNostrum, that are not equipped with GPUs.

The key element of the deployed layered architecture is Apache Spark (Zaharia et al.). In order to isolate machine-learning applications from the particularities of MareNostrum, Spark is usually used as an intermediate layer (not only in MareNostrum, Wang et al. (2014) does the same on a Cray X-series supercomputer). The deployment of Spark-enabled clusters over MareNostrum is not trivial, it has required the development of a specific interoperability layer that we call Spark4MN, which will be explained later. On top of this stack (Marenostrum, Spark4MN and Spark) we place a deep learning specific layer, DL4J. DL4J, that is written in Java and has a direct integration with Spark, enables distributed training of deep neural networks through a synchronous data parallelism method.

These four elements (DL4J, Spark, Spark4MN and MareNostrum) have been integrated enabling to efficiently train deep neural networks. Apart from the deployment details, the challenge is scalability and proper configuration. Simply running on many cores may yield poor benefits or even degraded performance due to overheads. We deal with this issue and we aim to make the first step towards systematic analysis of the several parameters and optimized configuration.

In order to evaluate the performance and scalability of

** Corresponding author: Tel.: +34-93 4054044;
e-mail: `rtous@ac.upc.edu` (Ruben Tous)

the proposed software stack on MareNostrum, we have experimented with different workloads and different deployment setups (number of nodes, parallelism configuration, etc.). Through the following sections we explain the different components of the deployment in more detail. Then, we discuss the performed experiments and the obtained results, aiming to shed light onto the parameters that have the biggest impact and their effective configuration. We provide insights into how the job configuration on a traditional HPC setup can be optimized to efficiently run this kind of workloads. The derived conclusions should be useful to guide similarly complex deployments in the future.

2. Related Work

Several works have addressed the execution of deep learning workloads on large specific purpose clusters usually involving nodes equipped with GPUs. In Kurth et al. (2017), authors present a Caffe-based approach to execute deep learning workloads on a contemporary HPC system equipped with Xeon-Phi nodes. They use the Intel distribution of Caffe, that improves Caffe performance when running on CPUs. Authors report to be able, due to a hybrid approach, to overcome the limitations of synchronous systems scaling the training of a model up to thousands of nodes. In You et al. (2017), authors describe another method (tested over KNL clusters and multi-GPU clusters) with very good weak scaling efficiency (e.g. 92% for GoogleNet on 2176 cores with respect to a Intel Caffe baseline). Alternatively, distributed DNNs training can be deployed through an integrated software stack. Despite of the potential performance limitations, the possibility to take profit of thousands of under-utilized cores to alleviate the pressing demand of computational resources to train deep learning models with a solution with minimum cost and setup time is an option for many general-purpose HPC infrastructures, specially if they already provide the lower components of the stack. A common case are infraestructures with an Apache Spark abstraction layer. Enabling distributed DNNs training in these situations is straightforward through the integration of a DL4J layer. While the performance of Spark on HPC setups have been already studied by many works (e.g. Michael et al. (2014), Wang et al. (2014), Haut et al. (2017)), as far as we know, there are no previous works evaluating the feasibility and scalabity of a Spark-DL4J integrated solution when applied to an HPC setup. One potential limitation of DL4J is that its currently constrained to synchronous SGD-based training. The work described in Keuper and Pfreundt (2016) analyzes the main bottlenecks of the synchronous approach. The authors conclude that the issue is quickly turning into a vastly communication bound problem which is severely limiting the scalability in most practical scenarios.

3. Deep Neural Networks

Deep neural networks (DNNs) are layered compositional models that enable learning representations of data with multiple levels of abstraction. State-of-the-art DNNs include many variants, specialized in different domains (convolutional deep neural networks, recurrent neural networks, etc.). DNNs are usually trained by using iterative, gradient-based optimizers (typically mini-batch SGD) that drive a non-convex cost function to a local minima. In every iteration step, we use information about the gradient ∇E at the current point. In iteration step $[t+1]$ the weight update $\Delta w[t]$ is determined by taking a step (γ is the learning rate) into the direction of the negative gradient at position $w[t]$ such that (in the case of stochastic training):

$$\Delta w[t] = -\gamma \frac{\partial E_n}{\partial w[t]} \qquad (1)$$

State-of-the-art networks have a huge number of weights W and the core computation in their training is dominated by dense linear algebra. Usually, in order to improve the efficiency, the training dataset is split into mini-batches of size B (typically chosen between 1 and a few hundreds) and the model is only updated (one iteration) after accumulating the gradients of all the training samples within a mini-batch.

DNNs training on a single node involves several software and hardware layers. At the top of the stack there is normally a deep learning framework such as DL4J, TensorFlow, Torch, etc. (there may be even an upper layer such as Keras). Below, the framework relies on an underlying numerical library such as NVIDIA's cuDNN or Intel's MKL. Finally, the models are usually trained on NVIDIA GPUs or Intel's Xeon Phi processors.

When trained on multiple nodes, one can apply data parallelism (distributing training samples among nodes) and/or model parallelism (distributing model parameters among nodes). In our deployment, we only apply data parallelism. The B training samples within a min-batch are split into n equal sized sets of size b (with $b = B/n$). The resulting mini-batch-splits are then fed to n nodes holding a complete copy of the model. The results (gradients) off all nodes are then accumulated and used to update the model.

While DL4J limits us to perform this process synchronously (awaiting all the workers to finish before updating the model), it could be also performed asynchronously (allowing model updates with just a part of nodes results). Asynchronous data parallelism can potentially gain higher throughput, but depending on the infrastructure status we can have the *stale gradient problem*. By the time a slow worker has finished its calculations based on a given state of the model, the model may have been updated a number of times and the outdated update may have a negative impact. Some solutions to this problem (e.g. Nguyen et al. (2018)) have been recently proposed.

4. DL4J

DL4J (or Deeplearning4j) is a computing framework written for Java with wide support for deep learning algorithms. DL4J is powered by its own numerical computing library, ND4J, and provides distributed parallel versions (both for GPUs and CPUs) of the algorithms that integrate with Apache Hadoop and Spark. Through a C++ native library, Libnd4j (with a BLAS backend), ND4J provides intra-node parallelism for matrix operations (implemented with OpenMP vectorizable loops with
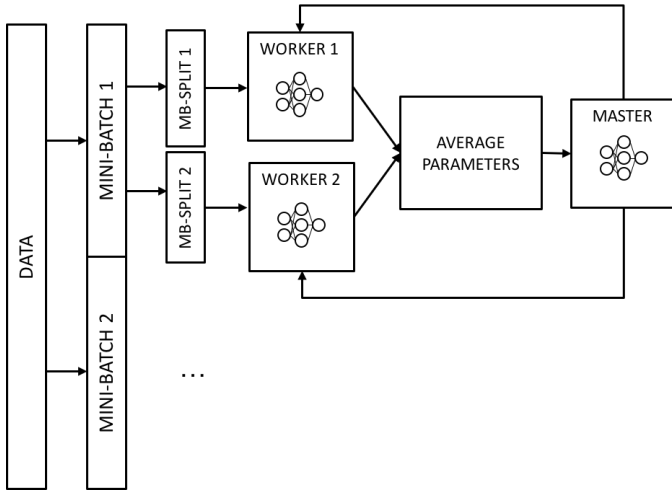
Fig. 1. Parameter averaging in DL4J over Spark (example using two mini-batches).

SIMD support). With the help of JavaCPP and JNI, pointers to off-heap memory (allocated outside of the JVM and not managed by the GC) are passed to the underlying C++ code. In order to achieve distributed network training over Spark, DL4J performs a version of the synchronous data parallelism mechanism called parameter averaging. Instead of transferring gradients to the master, the nodes perform first a local model update and then they transfer the resulting weights to the master, where they are averaged. With respect to generic parameter averaging, in DL4J the Spark driver and reduction operations take the place of the parameter server (see Figure 1).

There are several parameters that must be adjusted to optimize training time. These include, but are not limited to, mini-batch-split size, averaging frequency (too low averaging periods may imply too networking overhead), prefetching (how many mini-batch-splits a worker must prefetch to avoid waiting for the data to be loaded), and repartitioning strategy (when and how to repartition data to keep the partitions balanced).

5. Apache Spark

As mentioned before, Apache Spark is the key component of the proposed framework. Spark is a distributed system for processing data-intensive workloads. It excels in an efficient memory usage, outperforming Hadoop for many applications (Zaharia et al. (2012)). Spark is being used to execute big data workloads on the MareNostrum supercomputer, isolating the applications from the particularities of this HPC infrastructure. Spark is designed to avoid the file system as much as possible, retaining most data resident in distributed memory across phases in the same job. Such memory-resident feature stands to benefit many applications, such as machine learning or clustering, that require extensive reuse of results across multiple iterations. Essentially, Spark is an implementation of the so-called Resilient Distributed Dataset (RDD) abstraction (Zaharia et al. (2012)), which hides the details of distribution and fault-tolerance for large collections of items. The usage of Spark over alternatives with potentially better performance, e.g. MPI,

was a prerequisite of the described deployment. While Spark has some advantages such as fault tolerance, its main advantage over other alternatives in this case was that it minimized the deployment cost.

6. The Spark4MN Framework

The MareNostrum supercomputer is accessed through an IBM LSF Platform workload manager. In order to be able to deploy Spark clusters over MareNostrum, we employ an intermediate layer called Spark4MN (Tous et al. (2015)). Spark4MN is also in charge of managing the deployment of any additional resource Spark needs, such as a service-based distributed file system (DFS) like HDFS. Essentially, Spark4MN is a collection of *bash* scripts that deploy the Spark cluster's services, and executes the user applications. Spark4MN scripts read a configuration file, describing the application and the Spark cluster configuration, and submit one or more jobs to the MareNostrum workload manager. Once the cluster's job scheduler chooses a Spark4MN job to be executed, an exclusive number of cluster's nodes are reserved for the Spark cluster and (if requested) for the DFS (e.g. HDFS) cluster (may be the same nodes, depending on the configuration). After the resource allocation procedure, Spark4MN starts the different services. In Spark4MN, the Spark master corresponds to the *standalone* Spark manager, and workers are Spark worker services, where the Spark executors are received and launched. The cluster startup requires about 12 seconds. This is independent of the size of the cluster (the number of nodes). Each application is executed via *spark-submit* calls. During each Spark job execution, intermediate data is produced, e.g., due to shuffling. Such data are stored on the local disks and not on DFS by default (as in Michael et al. (2014), this yields the best performance). Finally, Spark timeouts are automatically configured to the maximum duration of the job, as set by the user.

7. Marenostrum supercomputer

MareNostrum is the Spanish Tier-0 supercomputer provided by BSC. It is an IBM System X iDataplex based on Intel Sandy Bridge EP processors at 2.6 GHz (two 8-core Intel Xeon processors E5-2670 per machine), 2 GB/core (32 GB/node) and around 500 GB of local disk (IBM 500 GB 7.2K 6Gbps NL SATA 3.5). Currently the supercomputer consists of 48896 Intel Sandy Bridge cores in 3056 JS21 nodes, with more than 104.6 TB of main memory and 2 PB of GPFS (General Parallel File System) disk storage. More specifically, GPFS provides 1.9 PB for user data storage, 33.5 TB for metadata storage (inodes and internal filesystem data) and total aggregated performance of 15GB/s. The GPFS filesystems are configured and optimized to be mounted on 3000 nodes. All compute nodes are interconnected through an Infiniband FDR10 network, with a non-blocking fat tree network topology. In addition to the 40 Gb/s Infiniband, 1 Gb/s full duplex Ethernet is in place. With the last upgrade, MareNostrum has a peak performance of 1.1 Petaflops.

8. Experiments and Results

The main goal of the experiments is to evaluate the scalability properties of the proposed deployment. To this end, we have experimented with different workloads and different deployment setups. Regarding the benchmarking workloads, we have chosen two widely used convolutional networks, AlexNet (Krizhevsky et al. (2017)) and GoogLeNet (Szegedy et al. (2015)). Both networks have been used in other state-of-the-art works (e.g. Keuper and Pfreundt (2016)) and let us compare our results with others. While AlexNet implements a rather shallow network with many parameters, GoogLeNet is a very deep network with many convolutional layers. We apply both networks to dataset of the ImageNet (Russakovsky et al. (2015)) visual recognition challenge. For reproducibility, we stick to the ILSVRC2012 classification task training and test datasets and their standard evaluation procedure.

Table 1. Properties of the deep neural networks used in the experiments.

	AlexNet	GoogLeNet
Default batch size	256	32
Default step-size	0.1	0.1
# Iterations till convergence	450k	1000k
# Layers	25	159
# Convolutional layers	5	59
# Fully-connected (FC) layers	3	1
# Weights in FC layers	55M	1M

Regarding the deployment setup, we have tested different values for the number of nodes, the number of Spark workers per node, the Spark data partition size, the DL4J averaging frequency and the persistence level. Figure 2 shows the speedup results obtained with $B = 256$ and $B = 1024$ (two Spark workers per node, averaging each 3 mini-batch-splits, Spark's persistence level set to *MEMORY_AND_DISK_SER* and automatic partitioning). The step sizes were increased according to the batch size as suggested by Iandola et al. (2016), while the number of iterations has been decreased by the same factor. For each different number of nodes n, each node processes mini-batch splits of size $b = B/n$.

Under a basic setup (averaging for each computed mini-batch and uniform node workload) synchronous data parallelism trough parameter averaging is mathematically equivalent to a non-parallel computation and yields the same accuracy results. However, accuracy degrades (regardless of parallelization) when mini-batch sizes become too large (Keskar et al. (2016)), which imposes a constraint on scalability. Table 2 shows the accuracy results for the different configurations.

Table 2. Impact on accuracy of the different configurations.

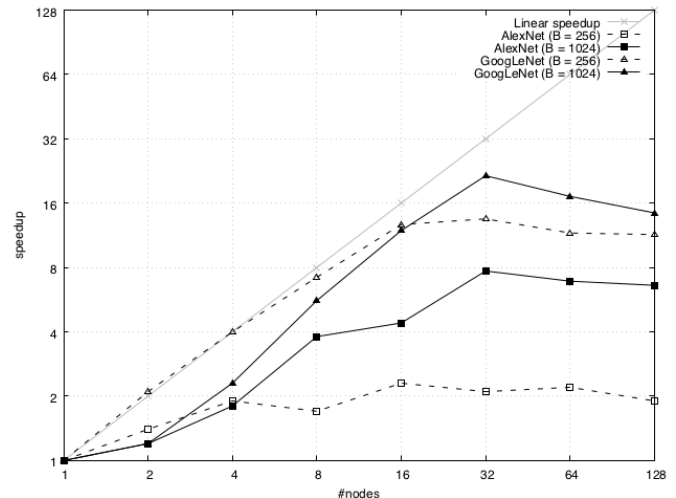| mini-batch size	accuracy
AlexNet (B = 256)	56.9%
AlexNet (B = 1024)	53.6%
GoogLeNet (B = 256)	67.1%
GoogLeNet (B = 1024)	65.4%

Fig. 2. Speedup results for AlexNet (Krizhevsky et al. (2017)) and GoogLeNet (Szegedy et al. (2015)) with different mini-batch sizes B.

The results of our evaluation show that DL4J and Spark are able to scale deep learning workloads over MareNostrum. However, the effective scaling stops above 32 nodes with the best configurations. This limitation agrees with the results reported in Keuper and Pfreundt (2016), that studies the theoretic constraints of synchronous data parallelism for DNNs training. The main bottleneck of the synchronous approach is the computation to communication ratio. The synchronous parallelization of DNN training requires the communication of the model w_t and the computed gradients Δw_t between all nodes in every iteration t. Since w has to be synchronous in all nodes and Δw_{t+1} can not be computed before w_t is available, the entire communication has to be completed before the next iteration $t + 1$. The problem is that w and Δw have the size of all weights in the neural network, which can be hundreds of megabytes. The compute times per iteration are rather low and decrease when scaling to more nodes. Depending on the model size and layout, the training problem becomes communication bound after scaling to only few nodes. Shallow networks with many neurons per layer (like AlexNet) scale worse than deep networks with less neurons (like GoogLeNet) where longer compute times meet smaller model sizes.

A second problem of the synchronous approach is that nodes process mini-batch-splits instead of mini-batches, and the size b of these splits depends on the number of nodes n. If b is too small (less than 32 samples in our experiments), there will be a negative impact on the inner parallel computation (within the node), especially in the case of the FullyConnected (FC) layers. One solution would be to increase the mini-batch size in parallel to the number of nodes, but too large batch sizes have been shown to cause slowdown in convergence and degrade the generalization properties of the trained model (Kurth et al. (2017)). As larger mini-batch sizes enable lower absolute training times when many nodes are used, they are always preferred, with the limitation of accuracy degradation. While intra-node parallelism of fully connected layers improve with a larger mini-batch size, the poor scalability results observed

at low core counts can be attributed to a poor intra-node parallelism of the other layers (dropout, pooling and LRN). The observed behavior is consistent with the results from Keuper and Pfreundt (2016). Another aspect negatively impacting scalability at low core counts (few nodes) can be related to data loading. Our implementation uses asynchronous data prefetching (as it is the default DL4J behavior). The next mini-batch-splits are loaded in another thread of the worker while training is proceeding in the main thread. Under ideal circumstances (fast disk access and small mini-batch-split size), asynchronous prefetching implies negligible data loading delays (except on the first iteration). However, MareNostrum nodes are equipped with relatively slow local disks (IBM 500 GB 7.2K 6Gbps NL SATA 3.5), a circumstance that can turn data loading into a bottleneck when mini-batch-splits are too big (i.e. when large mini-batches are distributed among few nodes) and the network is shallow.

A third problem is stragglers. The duration of the iteration depends on the slowest node. This effect gets worse with scale.

Asynchronous parallelization, not possible with the current version of DL4J, would solve these problems but, as mentioned before, has the *stale gradient problem* (though our nodes are homogeneous and the impact would be low). Some recent works like Kurth et al. (2017) propose a hybrid approach in which synchronous parallelism just takes place within groups of nodes. Partial solutions to the *stale gradient problem*, e.g. Nguyen et al. (2018)), have also been proposed.

9. Conclusions

The research work presented in this paper explores the feasibility and efficiency of using Apache Spark and DL4J for deploying deep learning workloads over a real-world, petascale, HPC setup, such as MareNostrum. To this end, we have designed a layered architecture consisting in both, general-purpose (Spark and DL4J) and custom components (Spark4MN). We have evaluated the deployment by training AlexNet and GoogLeNet over the ImageNet dataset. We have tested different deployment setups (number of nodes, number of Spark workers per node, data partition size, mini-batch size, mini-batch-split size, averaging frequency, prefetching and repartitioning strategy).

We conclude that it is feasible to rely on Apache Spark to deploy deep learning workloads over a traditional HPC setup. This approach minimizes deployment costs and enables a systematic tuning of the different configuration parameters, both at application level and at infrastructure level. However, the effective scaling is strongly limited by the synchronous parallelism approach applied by the latest DL4J version. Problems such as the communication overhead, mini-batch-split size and stragglers degrade the scalability beyond 32 nodes. In order to overcome this limitation, it would be necessary to replace the synchronous mechanism by a hybrid approach in which synchronization just takes place within fixed-size node sets. Assessing the impact of certain aspects, such as a quantitative evaluation of the effects on performance of asynchronous data prefetching, deserves further investigation and will be carried out in future work.

References

Fujimoto, N., 2008. Faster matrix-vector multiplication on geforce 8800gtx., in: IPDPS, IEEE. pp. 1–8. URL: http://dblp.uni-trier.de/db/conf/ipps/ipdps2008.html#Fujimoto08.

Haut, J.M., Paoletti, M., Plaza, J., Plaza, A., 2017. Cloud implementation of the k-means algorithm for hyperspectral image analysis. The Journal of Supercomputing 73, 514–529. URL: https://doi.org/10.1007/s11227-016-1896-3, doi:10.1007/s11227-016-1896-3.

Iandola, F.N., Moskewicz, M.W., Ashraf, K., Keutzer, K., 2016. Firecaffe: Near-linear acceleration of deep neural network training on compute clusters, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016, pp. 2592–2600. URL: https://doi.org/10.1109/CVPR.2016.284, doi:10.1109/CVPR.2016.284.

Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P., 2016. On large-batch training for deep learning: Generalization gap and sharp minima. CoRR abs/1609.04836. URL: http://arxiv.org/abs/1609.04836, arXiv:1609.04836.

Keuper, J., Pfreundt, F., 2016. Distributed training of deep neural networks: Theoretical and practical limits of parallel scalability, in: 2nd Workshop on Machine Learning in HPC Environments, MLHPC@SC, Salt Lake City, UT, USA, November 14, 2016, pp. 19--26. URL: https://doi.org/10.1109/MLHPC.2016.006, doi:10.1109/MLHPC.2016.006.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2017. Imagenet classification with deep convolutional neural networks. Commun. ACM 60, 84--90. URL: http://doi.acm.org/10.1145/3065386, doi:10.1145/3065386.

Kurth, T., Zhang, J., Satish, N., Mitliagkas, I., Racah, E., Patwary, M.M.A., Malas, T.M., Sundaram, N., Bhimji, W., Smorkalov, M., Deslippe, J., Shiryaev, M., Sridharan, S., Prabhat, Dubey, P., 2017. Deep learning at 15pf: Supervised and semi-supervised classification for scientific data. CoRR abs/1708.05256. URL: http://arxiv.org/abs/1708.05256.

Lee, V.W., Kim, C., Chhugani, J., Deisher, M., Kim, D., Nguyen, A.D., Satish, N., Smelyanskiy, M., Chennupaty, S., Hammarlund, P., Singhal, R., Dubey, P., 2010. Debunking the 100x GPU vs. CPU myth: an evaluation of throughput computing on CPU and GPU, in: 37th International Symposium on Computer Architecture (ISCA 2010), June 19-23, 2010, Saint-Malo, France, pp. 451--460. URL: http://doi.acm.org/10.1145/1815961.1816021, doi:10.1145/1815961.1816021.

Michael, S., Thota, A., Henschel, R., 2014. Hpchadoop: A framework to run hadoop on cray x-series supercomputers, in: Cray USer Group (CUG).

Nguyen, L.M., Nguyen, P.H., van Dijk, M., Richtárik, P., Scheinberg, K., Takác, M., 2018. SGD and hogwild! convergence without the bounded gradients assumption, in: Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, pp. 3747--3755. URL: http://proceedings.mlr.press/v80/nguyen18c.html.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Li, F., 2015. Imagenet large scale visual recognition challenge. International Journal of Computer Vision 115, 211--252. URL: https://doi.org/10.1007/s11263-015-0816-y, doi:10.1007/s11263-015-0816-y.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E.,
    Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich,
    A., 2015.   Going deeper with convolutions, in: IEEE
    Conference on Computer Vision and Pattern Recognition,
    CVPR 2015, Boston, MA, USA, June 7-12, 2015, pp.
    1--9.   URL: https://doi.org/10.1109/CVPR.2015.7298594,
    doi:10.1109/CVPR.2015.7298594.
Tous, R., Gounaris, A., Tripiana, C., Torres, J., Girona,
    S., Ayguadé, E., Labarta, J., Becerra, Y., Carrera, D.,
    Valero, M., 2015.   Spark deployment and performance
    evaluation on the marenostrum supercomputer, in: 2015
    IEEE International Conference on Big Data, Big Data
    2015, Santa Clara, CA, USA, October 29 - November 1,
    2015, pp. 299--306.   URL: https://doi.org/10.1109/
    BigData.2015.7363768, doi:10.1109/BigData.2015.7363768.
Wang, Y., Goldstone, R., Yu, W., Wang, T., 2014.
    Characterization and optimization of memory-resident
    mapreduce on hpc systems, in: IPDPS, pp. 799 -- 808.
You, Y., Buluç, A., Demmel, J., 2017.   Scaling deep
    learning on GPU and knights landing clusters.   CoRR
    abs/1708.02983. URL: http://arxiv.org/abs/1708.02983.
Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J.,
    McCauly, M., Franklin, M.J., Shenker, S., Stoica, I.,
    2012.   Resilient distributed datasets: A fault-tolerant
    abstraction for in-memory cluster computing, in: NSDI,
    pp. 15--28.
Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S.,
    Stoica, I., .   Spark: Cluster computing with working
    sets.